

Docket No.: 1267.1024

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:

Peter Raymond BALL, et al.

Serial No. 09/914,575

Group Art Unit: Unassigned

Confirmation No. 9631

Filed: August 31, 2001

Examiner: Unassigned

For: OPTICAL NETWORKS

SUBSTITUTE SPECIFICATION

TITLE OF THE INVENTION

OPTICAL NETWORKS

BACKGROUND OF THE INVENTION

1. Field of the Invention

5 This invention relates to optical networks.

2. Description of the Related Art

Figure 1A of the accompanying drawings shows in block diagram form the basic components of a passive optical network (PON). A multiwavelength optical source 3, located in a central office 1, transmits light signals consisting of multiple discrete wavelengths $\lambda_1 \dots \lambda_N$ down an optical fibre 10 to a wavelength division multiplexer (WDM) 7, located in a remote node 5, which then distributes the signals to a set of optical network units (ONUs) 9, via separate fibres 11. The network is described as passive since the optical routing components (such as the WDM 7) cannot actively be controlled or tuned during their operational use.

20 The wavelength division multiplexer 7 may be one of a variety of types. An example of a simple multiplexer is a power-splitting star coupler which simply splits incoming light into all ports equally; it is the trivial case of wavelength division multiplexing, because no selection is made on the basis of wavelength, and consequently all wavelengths $\lambda_1 \dots \lambda_N$ are distributed to all ONUs 9, as illustrated in Figure 1B of the accompanying drawings. This arrangement is sometimes referred to as "broadcast-and-select", since each signal is broadcast to multiple ONUs 9, and each ONU 9 then selects only those signals intended for it.

35 Instead of such a power-splitting star coupler, a wavelength routing element, for example an arrayed waveguide grating (AWG), could be used. An AWG splits incoming light into spectral constituents, launching

them onto separate output fibres. In this way, with an appropriately-designed AWG, incoming light consisting of wavelengths $\lambda_1 \dots \lambda_N$ could be multiplexed into N separate branches each consisting of light of only one of those wavelengths, as illustrated in Figure 1C of the accompanying drawings. In this way, each ONU 9 would only receive signals intended for that ONU, and each output branch would receive all the incoming power for its designated wavelength, unlike the star coupler where there is a splitting of power. Note that the architecture in Figure 1C shows the case where there are the same number of ONUs 9 as there are wavelengths emitted from the source 3, but this is not necessary; for example an ONU 9 could receive more than one of the routed wavelengths.

Figure 2 of the accompanying drawings shows an example of a recently-proposed two-stage wavelength-routed PON architecture having a multiwavelength optical source 3 at the optical line termination (OLT) emitting discrete wavelengths $\lambda_{11} \dots \lambda_{MN}$ down fibre 10. In the illustrated architecture there is one coarse AWG 4 located in an exchange 2, and M remote nodes 5, each having a fine AWG 7. Each fine AWG 7 feeds N ONUs 9, so that there are a total of $M \times N$ ONUs 9.

The coarse AWG 4 is designed to direct multiple wavelengths down each branch 6, and these wavelengths are then separated by the fine AWGs 7 and directed individually to each ONU 9 via the branches 11. This is achieved by ensuring that the free spectral range of the coarse AWG 4 is equal to the spacing of N channels received by the branches 11.

For example, using the illustrated architecture of Figure 2, the coarse AWG 4 receives at its input all wavelengths $\lambda_{11} \dots \lambda_{MN}$ emitted from the source 3. It directs wavelengths $\lambda_{11} \dots \lambda_{1N}$ down the first branch 6 to the first remote node 5. The AWG 7 within the first

remote node then directs each of the N wavelengths λ_{11} ... λ_{1N} at its input individually to the N respective ONUs 9.

In this architecture, like that of Figure 1C, each
5 ONU receives only the wavelength assigned to it, and for each wavelength there is no splitting of power at the routing components 4, 7.

The multiwavelength optical source 3 may be a
single tunable laser located in the optical line
10 termination (OLT) of fibre 10, constantly retuning and transmitting a different wavelength in different time slots. This scheme uses WDM principally to improve the privacy in the network; there is no increase in
15 capacity over a single wavelength system as only one wavelength is transmitted every time slot. The downstream protocol is effectively the same as a time division multiplexed (TDM) single wavelength system.

The above-described architectures of Figures 1C and 2 are fixed wavelength systems, since a wavelength
20 is permanently assigned to each branch of the PON, effectively creating a number of independent single wavelength networks within the same PON. This type of scheme is simple to implement but does not allow the redistribution of bandwidth in response to fluctuations in demand. For example, if the n'th ONU 9, permanently
25 assigned wavelength λ_n , is idle for a long period of time, then that wavelength λ_n is being wasted since it cannot be re-allocated to another ONU 9.

Dynamic assignment schemes seek to allow more
30 flexible use of bandwidth by introducing tunability into the network. The most obvious way to provide downstream wavelength re-allocation is to have tunable filters in the ONUs 9, in a broadcast-and-select
architecture such as that of Figure 1B where each ONU 9
35 receives more than one wavelength. The ONUs 9 of the PON would tune to the wavelength assigned to it in

response to a signal from the central office 1.

There is, however, a major drawback with this approach to dynamic assignment of wavelengths, which is that the information about current bandwidth requirements is held at the central office 1, and is separated from the location of the tunable components in the ONUs 9. Therefore when a retuning is required, a signal needs to be sent from central office 1 to the appropriate ONU 9, and an acknowledgement returned, before data destined to that ONU 9 can be transmitted on the new wavelength. As retuning is normally done in response to the overloading of a wavelength channel, this lag causes a build up of traffic and consequent increase in delay on that channel.

The present applicant has considered employing more than one tunable laser at central office 1 (the "head end") in a fixed wavelength PON so as to achieve a certain degree of dynamic bandwidth assignment without the use of tunable filters in the ONUs 9. With multiple tunable lasers, transmission to the ONUs 9 could be shared between the lasers. Each laser could be assigned its own set of ONUs to which to transmit, and consequently when the load on a particular laser is increased, for example due to an increase in demand from a particular ONU 9, responsibility for transmission to that ONU could be transferred to another less loaded laser. The effect would be effectively to transfer the tunability in the network from the ONUs to the head end.

There would be a number of advantages in doing this. Firstly, the tuning would be done with tunable transmitters rather than filters, the former currently having a faster tuning speed. Secondly, all the protocol functions would be controlled at the head end. Consequently, having the tuning there would mean that there is no delay between the tuning becoming necessary

and it being implemented. This could stop traffic build up on an overloaded transmitter as discussed above. Thirdly, the system would be more robust; if the tuning is at the ONU 9 then either an
5 acknowledgement of successful retuning is required, resulting in further delay, or there is the risk of an error in retuning resulting in the loss of cells transmitted to the ONU 9 on the new wavelength. Fourthly, the more expensive, tunable components would
10 be placed at the head end, where only a few are required, rather than providing expensive tunable systems at each ONU; this would lead to a cost reduction.

There are still certain drawbacks, however, to
15 such a fixed-filter, tunable-laser approach. Firstly, cells could be addressed to more than one ONU 9. This means that bandwidth would be wasted when the network transmits broadcast or multicast traffic, because the cell needs to be replicated and retransmitted on the
20 wavelength of each destination ONU 9. In contrast, a system with tunable filters at the ONU could be configured so that all the ONUs 9 in a multicast group can be tuned to the same channel. Secondly, constant retuning of the lasers at the head end would be
25 required. Consequently, if the tuning time is non-negligible, then a loss of bandwidth would result.

It is therefore desirable to provide a multiwavelength, broadcast-and-select optical network which combines head end tuning with efficient
30 transmission of broadcast and multicast traffic.

SUMMARY OF THE INVENTION

According to an embodiment of a first aspect of the present invention there is provided an optical network comprising: a plurality of optical network
35 units; and optical source means connected and arranged to transmit light signals to each of said plurality of

optical network units; wherein the said optical source means are capable of transmitting light signals at one or more of a plurality of different wavelengths, at least one optical network unit being operable to accept more than one of the said wavelengths, and each wavelength of the said plurality being accepted by at least one of the said optical network units such that each such wavelength is accepted by a different subset of optical network units, the optical network further comprising control means operable to cause the said optical source means to transmit light signals at one or more selected such wavelengths corresponding to respective desired subsets of the said optical network units.

According to an embodiment of a second aspect of the present invention there is provided control circuitry for use in an optical network, which network comprises a plurality of optical network units and optical source means connected and arranged to transmit light signals to each of said plurality of optical network units, said optical source means being capable of transmitting light signals at one or more of a plurality of different wavelengths, at least one optical network unit being operable to accept more than one of the said wavelengths, and each wavelength of the said plurality being accepted by at least one of the said optical network units such that each such wavelength is accepted by a different subset of optical network units, the control circuitry being operable to cause the said optical source means to transmit light signals at one or more selected such wavelengths corresponding to respective desired subsets of the said optical network units.

According to an embodiment of a third aspect of the present invention there is provided a dynamic bandwidth assignment method for an optical network

comprising optical source means capable of transmitting light signals at one or more of a plurality of different wavelengths, each of the said wavelengths being accepted by a different subset of optical network units of the said network, in which method: light signals are transmitted by the said optical source means at one or more wavelengths, selected from the said plurality of wavelengths, corresponding to one or more desired subsets of optical network units, and, in response to a required bandwidth redistribution, the said one or more wavelengths at which light signals are transmitted by the said optical source means are changed to one or more different wavelengths, selected from the said plurality, which correspond to one or more different desired subsets of optical network units.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1A shows in block diagram form the basic components of a multiwavelength passive optical network;

Figure 1B shows the passive optical network of figure 1A employing a power-splitting star coupler;

Figure 1C shows the passive optical network of figure 1A employing an arrayed waveguide grating;

Figure 2 shows in block diagram form the basic components of a two-stage passive optical network;

Figure 3 shows the basic principle of a staggered filter optical network architecture embodying the present invention;

Figure 4 shows the passband of the filter in each optical network unit of Figure 3;

Figure 5 shows the Figure 3 optical network architecture when the load is unbalanced;

Figure 6 shows the optical network units served by each laser of the Figure 3 embodiment for various laser tunings;

Figure 7 shows the optical network units served by each laser in another embodiment of the present invention;

5 Figure 8 shows the optical network units served by each laser in a further embodiment of the present invention;

Figure 9 shows the optical network units served by each laser in a yet further embodiment of the present invention;

10 Figure 10 shows the passband of the filter in each optical network unit of the Figure 9 embodiment;

Figure 11 shows the optical network units served by each laser in a yet further embodiment of the present invention;

15 Figure 12 shows the passband of the filter in each optical network unit of the Figure 11 embodiment;

Figure 13 shows a physical implementation of a staggered filter system embodying the present invention;

20 Figure 14 illustrates the use of a wavelength division demultiplexer as a filter;

Figure 15 shows the passband of the filter in each optical network unit in another embodiment of the present invention;

25 Figure 16 illustrates an example of the required transmissions for a wavelength-routed system not embodying the present invention;

Figure 17 illustrates the required transmissions for the Figure 16 example for a staggered filter system embodying the present invention;

30 Figure 18 is a table showing the typical percentage of viewers watching various television channels at the peak viewing hour;

Figure 19 is a graph showing the penetration rates for cable television;

Figure 20 is a graph showing the hour-by-hour

average viewing figures;

Figure 21 is a graph showing estimated savings of bandwidth in an embodiment of the present invention;

Figure 22 is another graph showing estimated savings of bandwidth in an embodiment of the present invention;

Figure 23 shows an embodiment of the present invention applied to a ring network architecture;

Figure 24 shows an embodiment of the present invention applied to a bus network architecture;

Figure 25 shows an example of a unidirectional optical coupler for use in the Figure 23 and Figure 24 embodiments; and

Figures 26 and 27 show possible designs of a bidirectional optical coupler for use in the Figure 23 and Figure 24 embodiments.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 3 shows the basic principle of a staggered filter optical network architecture embodying the present invention. In this embodiment, there are two tunable lasers 3 capable of transmitting a total of 10 wavelengths λ_1 to λ_{10} to a total of eight optical network units 19, ONU_1 to ONU_8 . The first tunable laser 3 is capable of transmitting one of five wavelengths λ_1 to λ_5 , and in the present example is tuned to λ_3 . The second tunable laser is capable of transmitting one of five wavelengths λ_6 to λ_{10} , and in the present example is tuned to λ_8 . Control portion 18 is in communication with, and controls the operation of, the tunable lasers 3.

Each optical network unit 19 employs a bandpass filter which allows a group of five consecutive transmitted wavelengths to be passed. For example, as indicated in Figure 3, ONU_1 passes wavelengths λ_1 to λ_5 , while ONU_4 passes wavelengths λ_3 to λ_7 . The table of Figure 4 summarises the passband of the filters in each

of the ONUs 19, where shaded boxes indicate those wavelengths that are passed by the appropriate filter. It can be seen that, in this embodiment, the passband of neighbouring ONU filters are overlapping and form a staggered progression from one wavelength limit to the other. Each laser transmits to every optical network unit, but since each optical network unit filters out certain wavelengths, not every optical network unit will actually respond to signals of each wavelength.

It is also apparent from the table of Figure 4 that when the control portion 18 causes the first laser 3 to be tuned to λ_3 and the second laser 3 to be tuned to λ_8 , all of the ONUs 19 are served by one or other of the lasers, with ONU₁ to ONU₄ being served by the first laser and ONU₅ to ONU₈ being served by the second laser. This arrangement is suitable when there is a balanced load, with each laser serving the same number of ONUs. Should there be an increase in demand from ONU₁ to ONU₃, for example, it is possible with this architecture to perform a limited degree of bandwidth re-distribution by transferring ONU₄ from being served by the first laser to being served by the second laser, thereby freeing more time for the first laser to serve the more demanding ONU₁ to ONU₃.

This is done by making use of the staggered nature of the filters, and by coordinating the retuning of the two lasers. The control portion 18 causes the first laser to be retuned to λ_2 , and the second laser to be retuned to λ_7 , as shown in Figure 5. Since ONU₄ filters out the new first laser wavelength λ_2 , but passes the new second laser wavelength λ_7 , it has effectively been transferred from the first to the second laser. The first laser now transmits only to the first three ONUs 19.

In this way it can be seen that wavelengths are assigned to groups of ONUs 19 rather than uniquely to

one. The tunable lasers can target different groups of ONUs by using different wavelengths. By coordinating the retuning of the different lasers, the allocation of ONUs to head end lasers can be changed according to changes in demand distribution. A table showing the listing of the ONU groups and the required laser tunings is shown in Figure 6, from which it can be seen that tuning the lasers to λ_1, λ_6 (or λ_5, λ_{10}) produces a 2-6 (6-2) distribution and tuning the lasers to λ_2, λ_7 (or λ_4, λ_9) produces a 3-5 (5-3) distribution.

The staggered nature of the filters illustrated in Figure 4 allows a complete coverage of all of the ONUs 19 by transmission on either of two wavelengths. If those two wavelengths are chosen appropriately, then the coverage can be achieved without being able to transmit to any one ONU 19 on both wavelengths; for example λ_1 and λ_6, λ_4 and λ_9 , or λ_5 and λ_{10} . In this way the receiver in the ONU 19, which responds to light energy rather than to a particular wavelength, will not receive confusing signals.

There are a large number of ways in which the wavelengths and ONUs can be grouped to allow different re-distributions. There is a trade-off between the number of wavelengths used and the complexity of the filtering arrangements on the one hand and the possible degree of retuning and the consequent network benefits on the other.

The system may be altered to provide greater or less tuning according to the number of wavelengths that are used and the bandwidths of the filters. Figure 7 shows a table illustrating the use of fourteen wavelengths to produce a system that allows up to seven ONUs 19 to be assigned to one laser, and the table in Figure 8 shows the scenario of allowing a maximum of five ONUs 19 on each laser, which allows a reduction in the number of wavelengths to six. The filtering

requirements would, of course, change with the changing number of wavelengths. In the Figure 7 case, seven wavelengths would need to be passed by each filter, and in the Figure 8 case, only three.

5 Although the above architectures are straightforward in their physical implementation, in certain circumstances there may not be sufficient flexibility in the way that the re-distribution of ONUs 19 to the head end lasers can be done. The problem of
10 fairness arises in that all the ONUs 19 are not treated equally. For example, a surge in demand in one of the central ONUs (ONU₄ and ONU₅) cannot be compensated by a redistribution in the way that a similar increase in traffic from the "edge" ONUs can be.

15 In order to allow all the ONUs 19 to be re-distributed from one laser to another, it becomes necessary to allow certain ONUs 19 to receive wavelengths which are not consecutive in the ITU (International Telecommunication Union) grid.

20 The first example of this principle is a relatively simple one, which allows the re-distribution of any four (out of a group of 8 ONUs 19) consecutively numbered ONUs 19 to each laser using a total of 8 wavelengths. There are always the same number of ONUs
25 19 per laser, but a degree of flexibility in their distribution is allowed. Figure 9 shows the wavelength groupings and Figure 10 shows the filtering requirements. It can be seen from the table of Figure 10 that the filters in ONU₁ to ONU₃ pass both low and
30 high wavelengths.

 Although this arrangement is relatively straightforward and fair, it can be limited in its ability to provide dynamic bandwidth allocation in certain circumstances. In order to provide more
35 flexibility, a more complex method may be employed. An arrangement of 16 wavelengths that will allow a three-

five split between the lasers (i.e. one of the two lasers serves five of the 8 ONUs and the other laser serves the remaining three) is shown in Figures 11 and 12.

5 In order to provide a more complete re-allocation scheme, there can be the option of having four ONUs served by each laser or five on one and three on the other. This can be achieved by using both of the schemes described in Figures 9 to 12 in conjunction
10 with each other (i.e. using a total of 24 wavelengths).

 Greater flexibility, allowing for the provision of six-two splits, could also be achieved. However, the more re-allocation that is provided, the more wavelengths that are required and the more complex the receiver arrangements at the ONUs.
15

 Figure 13 shows a simple physical implementation of the staggered filter scheme. In this example, there are four tunable lasers 3 controlled by control portion 18, and the signals emitted therefrom are multiplexed
20 by multiplexers 12 for transmission down single fibre 10. A passive splitter 7 (such as a power-splitting star coupler described above in relation to Figure 1B) distributes the signals down branches 11 to the remote units. Appropriately-selected bandpass filters 13 pass
25 only the required wavelengths on to a receiver 14. Each filter/receiver pair may be located, for example, within an optical network unit. Optical amplifiers can also be used along increased optical fibre spans to compensate for any losses or attenuation.

30 The filtering can be implemented using fixed filters, manually tunable Fabry-Perot filters or slow tunable filters. The first option is the cheapest in terms of component cost, but may cause problems for the network operators if different components are required
35 for each ONU 19.

 The most basic form of the staggered filter

architecture can be implemented using bandpass filters. The requirements for the optical network detailed in Figure 4 are for filters with a bandwidth equivalent to 5 wavelengths on the ITU grid and a free spectral range of greater than 40 ITU wavelengths. Since filters with bandwidths between 0.25 and 100nm are commercially available, this does not pose a problem.

The more flexible architectures (such as those illustrated in Figures 9 to 12) cannot be implemented in such a straightforward manner, as the wavelengths they need to receive are not always consecutive in the ITU grid. One way to satisfy a more complicated filtering requirement would be to exploit the periodic nature of optical filters. The architecture in Figure 10, for example, could be implemented by arranging for λ_1 , λ_2 and λ_3 to be in a passband one free spectral range apart from λ_6 , λ_7 and λ_8 .

For the more complex systems, such as that outlined in Figure 12, wavelength division demultiplexers can be used to separate and select the incoming wavelengths. The WDM would be able to separate all the wavelengths present in a branch of the optical network, but only the ports carrying the required wavelengths for the particular ONU would be connected to the receiver. The arrangement required for ONU₁ in the Figure 12 architecture is shown in Figure 14. The selection of wavelengths can be done optically, for example by connecting only the required ports to the receiver, or electronically, by having, for example, a photodiode at each port and only sending the required signals to the receiver amplifiers.

It will be appreciated that, although embodiments of the present invention have been described which employ two or four tunable lasers at the head end, each tuned to a fixed wavelength until bandwidth reallocation is required, other embodiments are not

limited to this.

For example, in the Figure 3 embodiment a single tunable laser capable of emitting all wavelengths λ_1 to λ_{10} could be used instead of the two lasers. The network would then operate in a time division multiplexed manner, with a single wavelength occupying each time slot. In the Figure 3 example, the single tunable laser would constantly re-tune, transmitting λ_3 in one time slot and then λ_8 in another.

In addition, as demonstrated by Figure 13, other embodiments are not limited to transmission on only two wavelengths at a time. For example, three or more lasers at the head end could be used.

The filtering arrangement shown in Figure 15 requires the use of three tunable lasers at the head end emitting at three different wavelengths (or one/two tunable lasers operating in a time division multiplexed manner). For example, the three lasers could be tuned to $(\lambda_1, \lambda_4, \lambda_7)$ or $(\lambda_2, \lambda_5, \lambda_8)$ or $(\lambda_3, \lambda_6, \lambda_9)$ respectively to cover completely all ONUs.

The issue of multicast traffic has already been raised above, with respect to the wavelength-routed architecture. This type of network has a unique wavelength for each of the ONUs in the system. The wavelength routers at the remote nodes ensure that only data destined to a particular ONU is sent there. Ordinarily a problem would arise when the same data is required to be sent to a number of different ONUs, since then the information has to be replicated on each of the wavelengths of the target ONUs.

With the staggered filter architecture, on the other hand, if there are several ONUs in a wavelength group that require the same data, the data only needs to be transmitted once.

An example of this is shown in Figures 16 and 17. Figure 16 shows a data sequence and the ONUs to which

each of the cells in the sequence need to be sent. The required transmissions and the laser tunings needed are shown in Figure 16 for the wavelength-routed system and Figure 17 for the staggered filter architecture. It is
5 apparent from Figure 16 that there can be significant replication of data and a consequent wastage of time in the wavelength-routed system, compared to the staggered filter system transmissions shown in Figure 17.

It should be noted that there is potential for
10 further improvement if there is a large amount of multicast traffic. This would allow the assignment algorithm to take account of the multicast groups when assigning ONUs to the transmitters.

A quantitative estimate of the benefits of the
15 staggered filter system against the simple addition of extra tunable lasers to the wavelength-routed approach will now be calculated. The approach used here is to assume that there is a given probability of a user being a member of a multicast group. From this, the
20 probability of an ONU containing at least one member of that multicast group is calculated; such an ONU will be referred to as a member ONU. The expectation value of the number of member ONUs in a group served by one tunable laser is determined, and from this, the average
25 number of cell replications saved. The statistics used in the estimate are for Cable TV (CATV), working on the assumption that cells for a given TV channel are sent only to the ONUs where that channel is being watched. The figures are for channels watched at peak viewing
30 hours, based on historical data.

Figure 18 shows the percentage of viewers watching the most popular channels at the peak viewing hour (2000-2100). This information can be combined with the penetration rates for cable television (Figure 19) and
35 the hour by hour average viewing figures (Figure 20) to calculate the probability of a user on the optical

network watching a particular channel on CATV during the peak hour. The optical network configuration being considered here is fibre-to-the-cabinet (FTTCab) in which the ONUs are situated in a street cabinet.

5 Signals to and from customers are multiplexed at the ONU so that, for example, between 8 and 128 customers can be supported per ONU. The next stage is to calculate the probability of a given ONU requiring the channel in question and hence the expectation value of
10 the number of ONUs requiring the channel that are served by a single tunable laser. This figure allows an estimate of the benefits of the staggered filter architecture compared with the wavelength-routed system, as the latter would require a copy of the data
15 for each ONU whereas the former requires only one copy for the group.

The proportion of customers on the optical network with CATV is assumed to be approximately equal to the total number of CATV customers divided by the number of
20 households in the country (around 20 million). This gives a proportion of approximately 0.25. Figure 20 shows the percentage of customers watching CATV during a 24 hour period. Taking the peak viewing hour figure and combining this with the data shown in Figure 18
25 produces statistics for viewing of individual channels.

The bandwidth saving achieved by using the staggered filter system rather than the wavelength-routed architecture can be estimated by the following method.

30 The probability of any one customer on an ONU using a given channel is calculated by using a binomial distribution. This is given by $1 - \text{Prob}(\text{no customers watching a given channel})$.

35 With a wavelength-routed architecture, the data stream for a TV channel has to be replicated for every ONU requiring the service. For a staggered filter

system the data stream need only be transmitted once for each group where one of the ONUs requires the channel. Hence there is a saving equal to the channel bandwidth for each ONU after the first one to require the service in a given group.

It is assumed that the average number of ONUs in a group served by a single tunable laser is eight (four channels serving a standard 32 way split). The saving is then calculated as follows:

$$\text{Saving} = \frac{\sum_{i=2}^8 (i-1) \times \text{Prob}(i \text{ customers})}{\sum_{i=1}^8 i \times \text{Prob}(i \text{ customers})}$$

This is then summed over all the channels shown in Figure 18.

The estimated savings in bandwidth (in terms of multiples of the channel bandwidth) against the number of customers per ONU are shown in Figure 21 and range from approximately 30% for 8 customers per ONU up to approximately 75% for 128 customers per ONU. The savings in bandwidth against the number of ONUs per wavelength group are shown in Figure 22 for various numbers of customers per ONU and range from approximately 15% to 45% for 8 customers per ONU to approximately 55% to 85% for 128 customers per ONU.

The overall saving compared with the total capacity of the system can be calculated by estimating the proportion of the total bandwidth of the system that is used for cable television. Further demand predictions will be required before a conclusion can be fully drawn on the bandwidth savings of the scheme.

One of the features of the wavelength-routed architecture is the need for constant retuning of the laser in order to transmit to different ONUs. The

tuning time of the laser is therefore of critical importance. If the tuning latency approaches the transmission time of an ATM cell, then bandwidth will be lost, as time between transmissions will need to be dedicated to retuning. If tuning times are significantly greater than the length of an ATM cell, then the network becomes increasingly impractical.

By contrast, retunings are only needed in the staggered filter architecture in response to changes in the demand on the network, most of which will occur only at call setups.

Tuning times in the order of nanoseconds have been reported in the literature, but commercially-available devices lag behind this. For example, the GEC four section laser is a prototype, expected to be commercially available within the next year. There are two effects which determined the tuning time, namely electronic and thermal. The electronic effects are in the order of 1-10 ns and the thermal effects in the order of ms. The relationship between current and wavelength change is not however linear. Four different currents are used to tune the laser, namely, the gain, coupler, phase and reflector. Work is currently underway to develop algorithms to minimise the changes in each current for any required change in wavelength. In this way, the thermal effects can be reduced, thus reducing the overall tuning times. A tuning time of 500 ns may be possible. Such a time would indicate a gain in bandwidth of approximately 50% for the staggered filter system over the wavelength-routed architecture.

The above-described arrangements use the staggered filter design with a passive optical network architecture like that shown in Figure 13. The staggered filter concept can also be applied to other network arrangements (which may contain active

components), for example a ring architecture and a bus architecture where the traffic distribution is all to/from one hub node, as shown in Figures 23 and 24 respectively. In these Figures, Node 1 is the hub node. It transmits signals at selected wavelengths using a tunable laser source similar to the method described above. Nodes 2 to n in these Figures are designed to access selected wavelengths using the staggered filter approach described above.

The choice of the number of wavelengths available, the number of nodes served and the wavelength selection at each node is the same as that described above.

As for the above-described architectures, the staggered filter arrangement is for 'downstream' traffic only. Upstream traffic can be supported using a separate fibre or the same fibre by using a different wavelength(s).

Signals are coupled in and out of the ring or bus using an optical coupler. Designs for the optical coupler are given in Figures 25 to 27. Figure 25 shows a design for use in a unidirectional ring or bus (traffic only transmitted in one direction on the fibre). Figures 26 and 27 show designs for a bi-directional ring or bus (traffic transmitted in both directions over a single fibre). It is equally possible to use one fibre for each direction of transmission.

The terminal equipment may use any form of optical transmission (e.g. including Synchronous Digital Hierarchy (SDH) or Plesiochronous Digital Hierarchy (PDH)).